

Express Mail Label No.: EL659733909US

Date Mailed: April 4, 2001

UNITED STATES PATENT APPLICATION FOR GRANT OF LETTERS PATENT

**Dexter Chun
Kou-Chun Lee
Robert Knight
Ravi Palakodery
INVENTOR(S)**

ROBUST RADIO BASE STATION CONTROLLER ARCHITECTURE

**COATS & BENNETT, P.L.L.C.
P.O. Box 5
Raleigh, NC 27602
(919) 854-1844**

ROBUST RADIO BASE STATION CONTROLLER ARCHITECTURE

BACKGROUND OF THE INVENTION

The present invention relates to wireless communication systems in general, and

5 particularly relates to base station architectures for next generation CDMA radio access networks.

Most first and second-generation CDMA radio access networks use proprietary network architectures and protocols to provide interconnection between the core of the radio access networks (e.g. MSC) and the base transceiver stations (BTSs). One such proprietary network is known as the base communication network (BCN). A BCN employs a star-cluster topology in which clusters of processing devices are attached to each leg of the star. Base station controller (BSC) functionality is distributed among these processing devices. Some level of robustness is achieved in the BCN architecture by distributing mission critical functions on different legs of the star, and by using a redundant hub that constantly monitors its own health.

Demand for wireless services has increased dramatically in recent years. This increased demand has exposed a number of limitations in the BCN architecture. The BCN links are limited in bandwidth to about 10 Mbps. The BCN transport bandwidth places a ceiling on the number of Erlangs that can be supported and the physical

20 volume of equipment required to increase call volume further makes it impractical to evolve the architecture. In addition, because multiple devices share the same BNC links, which are not fault tolerant, a transport failure can result in a significant number of dropped calls. Also, a proprietary communication protocol requires the use of special analyzers to observe information flow and monitor performance of the radio access

25 network.

BRIEF SUMMARY OF THE INVENTION

A base station controller platform comprises a redundant, distributed switching fabric flexibly interconnecting needed call processing resources in a fault tolerant fashion. Preferably, the switching fabric comprises a high data rate transport based on ATM. A central ATM-based switching resource interfaces with the different call processing resource pools via a number of distributed ATM switches. With this approach, call-processing flows may be established through the base station controller using the specific resources needed for the particular type of call, based on configuring the switching fabric for the appropriate resource interconnection.

10 The disclosed architecture provides high call density and scalability based on the high data rate capacity of the switching fabric and the arrangement of resource pools. In a rack arrangement, a hub subrack carries a centralized switching resource to support interconnection with one or more processing subracks. The processing subracks may be configured to each carry the different types of processing resources needed to

15 support communication between associated radio base stations and, for example, a mobile switching center or a packet data network. In this manner, a system controller can essentially establish a complete call flow through the base station controller on a single processing subrack. Performance is improved by minimizing the need to pass call-related data across processing subracks via the centralized switching resource.

20 Of course, the flexibility of the distributed, redundant switching fabric permits other rack architectures. For example, the different types of processing resources may be organized by processing subrack, with each type of processing subrack independently accessible via the centralized switching resource. In any case, the disclosed base station controller architecture provides for high call density. For

25 example, the disclosed architecture provides call capacity in excess of 6000 Erlangs within a volume no larger than three industry standard equipment cabinets. This high

call density is bolstered by the fault tolerance afforded by the redundant transport (communication) links interconnecting the call processing resources, which provides high equipment availability ratings necessary to meet applicable reliability requirements.

5

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a diagram of a prior base station control system.

Fig. 2 is a diagram of a general architecture for a base station control system in accordance with the present invention.

Fig. 3 is a more detailed diagram of the base station control system of Fig. 2.

10 Fig. 4 is a diagram of an exemplary subrack arrangement for the base station control system of Fig. 3.

Fig. 5 is a diagram of an alternate exemplary arrangement for the base station control system of Fig. 3.

15 Fig. 6 is a diagram of the distributed switching fabric used in the base station control system of the present invention.

Fig. 7 is a diagram of flow logic for base station control system call processing resource allocation for the subrack architecture of Fig. 4.

DETAILED DESCRIPTION OF THE INVENTION

20 The present invention is described in terms of an advanced base station control system (BSC) for use in a third generation (3G) Code Division Multiple Access (CDMA) wireless communication network. Some of the nomenclature used in describing the various call processing resources is specific to the assignee of the instant application, but the resource functionality will be readily apparent to those skilled in the art.

25 Moreover, it should be understood that the BSC architecture described below may, due

to its novel interconnection and flexible call processing flow, be implemented in a number of different ways.

Turning now to the drawings, Fig. 1 is a diagram of a wireless communication network 10 comprising a typical CDMA BSC 12, one or more mobile switching centers 5 14, one or more radio base stations (RBSs) 16, and a plurality of wireless access terminals 18, also referred to as user terminals (UT) 18. The BSC 12 comprises a plurality of selector bank subsystems (SBSs) 20 providing communication interface and processing functions in support of voice and data calls to and from the UTs 18. A CDMA Interconnect System (CIS) 22 provides interconnection between the SBSs 20 and a plurality of RBSs 16. The SBSs 20 may interface to the Internet 32 through a router 10 switch 21. A base station manager or controller 24 provides overall control of the BSC 12, and a time/frequency unit (TFU) 26 provides timing reference signals for coordination and synchronization of the BSC 12 within the network 10.

In operation, the BSC 12 cooperates with the MSC 14 to provide communication 15 between the UTs 18 and the PSTN 30 and, possibly, the Internet 32. Depending upon the type of call placed from or to a given one of the UTs 18, the BSC 12 performs certain signal processing and call management functions. For example, for a voice call, the call is routed through one of the SBSs 20, where voice encoding and decoding (vocoding) is performed, along with echo cancellation processing and certain radio link management 20 operations. Call data from an SBS 20 is packetized for transfer to the appropriate RBS 16 via the CIS 22, which operates as a relatively high capacity switch, passing call data to and from the RBSs 16 and SBSs 20.

Each SBS 20 includes a portion of the overall call processing and support 25 resources contained within the BSC 12. A certain number of the communication links (e.g., T1/E1 lines) between the BSC 12 and MSC 14 are routed into each of the SBSs 20. Communication traffic does not cross between SBSs 20 and call processing

resources are not shared between SBSs 20. Thus, the loss of a single SBS 20 results in the loss of a potentially significant portion of the overall call processing capacity of the BSC 12. Moreover, the arrangement and design of SBSs 20 is such that the processing or functional flow of a given call is substantially constrained to the flow defined by the 5 architecture of the SBS 20. That is, the processing path of a given call is determined first by the SBS 20 to which it is allocated, and by the architecture of the SBS 20 itself.

Fig. 2 is a simplified diagram of a base station controller (BSC) 40 in accordance with the present invention. The BSC 40 comprises a distributed switching fabric 42, a number of communication processing resource pools 44, one or more general-purpose processor boards (GPBs) or controllers 46, and a Timing Unit Board (TUB) 48. 10 Typically, the BSC 40 includes redundant TUBs 48. The BSC 40 interfaces with one or more MSCs 14, a packet data network serving node (PDSN) 50, and one or more RBSs 52. Generally, the BSC 40 communicatively couples a plurality of UTs 18 a core network that typically comprises the MSC 14 and the PDSN 50.

15 In general, each resource pool 44 is an overall collection of like processing devices or resources that provide one or more aspects of call processing, control, and management function within the BSC 40. The arrangement of switching fabric 42 and resource pools 44 provides a number of significant advantages. For example, the BSC 40 may be made exceptionally fault tolerant by making the switching fabric 42 redundant, such that resource pools 44 are accessible via two or more communication 20 links. Further, the call processing flexibility of the BSC 40 is significantly enhanced because the particular functional flow for a given call or type of call is determined by the configuration of the switching fabric 42 under control of the controller 46. That is, the controller 46 chooses the specific processing resources allocated to a given call by 25 configuring the switching fabric 42 to establish a given set of interconnections dedicated to that call through the BSC 40. Additionally, the scalability of the BSC 40 is significantly

improved as the switching fabric 42 flexibly supports the addition of new or expanded resource pools 44.

Fig. 3 is a more detailed diagram of the BSC 40 and illustrates two typical call processing flows, one for voice and traditional data/fax calls, and one for packet data

5 calls. The BSC 40 typically includes one or more of the following resource pools 44:

- front haul exchange termination (FH-ET) resources 44A providing termination for communication links between the BSC 40 and one or more MSCs 14;
- service option element (SOE) resources 44B providing selected signal processing functions such as vocoding and echo cancellation for voice calls, and pass-through functions for traditional asynchronous data/fax calls;
- selector element (SE) resources 44C providing radio link management such as handover and outer loop power control and signaling functions, and for packet data calls the SE resources 44C provide protocol support for TCP/IP, and for async data/fax they provide the radio link protocol (RLP) and the intersystem link protocol (ISLP);
- back haul exchange termination (BH-ET) resources 44D providing termination between the BSC 40 and one or more RBSs 52; and
- packet core network exchange termination (PCN-ET) resources 44E providing termination between the BSC 40 and a packet data serving node (PDSN)

20
50.

The switching fabric 42 preferably comprises a number of Asynchronous Transfer Mode (ATM) switches, or possibly a number of high data rate Ethernet switches. ATM switching includes the advantages of data rates in excess of 600 Mbps 25 and the ATM switching protocol's independence of data type. That is, with ATM, call traffic and management information passes through the switching fabric 42 in the same

53-byte "cells" regardless of whether the call is a voice, async data/fax, or packet data call. Also, the ATM layer is independent of the type of physical link between resource pools 44. Thus, the links between the resource pools 44 and the switching fabric 42 may be electrical or optical as needed or desired without changing the underlying
5 protocol.

ATM implementation of the switching fabric 42 has the further advantage of straightforward assignment and management of quality of service (QoS) for calls routed through the BSC 40. QoS may be easily defined for individual calls or groups of calls, and supported by appropriate channel assignments within the ATM-based switching
10 fabric 42.

The switching fabric 42 may also be implemented as an IP-based switch. In that type of implementation, the switching fabric 42 becomes a packet-switched arrangement, rather than the circuit-switched arrangement associated with the ATM-based implementation. In either case, the available devices or resources within each
15 resource pool 44 are accessible for call processing as needed through the switching fabric 42.

As noted, Fig. 3 illustrates two typical call processing flows, a first functional path supporting the more traditional voice and async data/fax calls to and from the UTs 18, and a second functional path supporting packet data calls between the UTs 18 and a
20 packet data network (PDN) 32, such as the Internet. The controller 46 configures each of these functional paths based on configuring the switching fabric 42 to access specific resources within each of the resource pools 44 involved in supporting the functional flow.

For voice and async data/fax calls, the processing route from the MSC 14 to a given one of the access terminals (UT) 18 is:

25 PSTN→MSC→FH-ET→SOE→SE→BH-ET→RBS→UT.

Where the italicized elements represent resources within the BSC 40. From the UT 18, the call processing route is:

UT→RBS→*BH-ET*→SE→*SOE*→*FH-ET*→MSC. Note also that similar functional flows may be established for communication between UTs 18.

5 For packet data calls, the processing route from the PDN 32 to a given one of the UTs 18 is:

PDN→PSDN→*PCN-ET*→SE→*BH-ET*→RBS→UT.

From the UT 18 to the PDSN 50, the call processing route is:

UT→RBS→*BH-ET*→SE→*PCN-ET*→PDSN→PDN.

10 Note that with CDMA soft handoff, calls may be simultaneously routed through multiple RBSs 52, and this is implicit in the above call routing flows.

As noted above, the controller 46 configures the switching fabric 42 to select specific processing resources, such as a particular digital signal processor (DSP)—not shown—from within the SOE resource pool 44B, by configuring the switching fabric 42. 15 More specifically, in making a SOE resource assignment for a given call, the controller 46 selects an available SOE resource, such as a DSP, from the SOE resource pool 44B, by configuring the switching fabric 42 to access the specifically assigned DSP from the SOE resource pool 44B.

Thus, the individual resources within the various processing pools 44 are 20 essentially available on a global basis according to the needs of a given call routed through the BSC 40. Configuration of the switching fabric 42 to support the functional flow (processing path) for a given call creates a dedicated set of circuit resources and transport channels for that call. The BSC 40 releases these dedicated resources upon termination of the call.

25 Fig. 4 is a diagram of an exemplary rack/subrack architecture for the BSC 40. With the “mixed” subrack architecture depicted, the fault tolerance and scalability of the

BSC 40 are enhanced. These enhancements yield direct benefits to service providers seeking maximum call processing density with minimal loss of call processing capability arising from single-point equipment faults.

In this mixed architecture arrangement, the BSC 40 comprises a "hub" subrack

5 60 and one or more processing subracks 62. In the illustrated example, the BSC 40 comprises processing subracks 62A through 62N. Each processing subrack 62 includes a portion of resources from each type of resource pool 44. That is, each processing subrack 62 includes a portion of the FH-ET resource pool 44A, a portion of the SOE resource pool 44B, a portion of the SE resource pool 44C, and a portion of the BH-ET resource pool 44D. In this arrangement, a particular resource pool 44, the SOE resource pool 44B for example, comprises the overall set of SOE resources distributed across the number of processing subracks 62A ... 62N installed in the BSC 40.

The hub subrack 60 comprises primary and secondary central switching cores that interface with primary and secondary processing subrack switches, one or more GPBs (controllers) 46, one or more TUBs 48, and the PCN-ET resource pool 44E. Together, the hub subrack and processing subrack switching resources comprise a primary distributed switching fabric 42A and a secondary distributed switching fabric 42B. The switching fabric 42 provides access between GPBs 46 and the resource pools 44, including the PCN-ET 44E resources located on the hub subrack 60.

20 Communication links between the hub subrack 60 and the various processing subracks 62 comprise redundant primary links 45A and secondary links 45B. Thus, the controller 46 may assign specific processing resources in any one of the processing subracks 62 to support a given call routed through the BSC 40, independent of the other resource assignments made for that call. For example, assume that a given call 25 originates via the PSTN 30 and terminates through the MSC 14 into the portion of the FH-ET resource pool 44A carried by the processing subrack 62A. In support of the call,

the controller 46 may assign SOE resources 44B from the processing subrack 62B, and SE resources 44C and BH-ET resources 44D from the processing subrack 62N, depending upon the availability of specific processing resources in the overall set of subracks 62. This flexibility in resource allocation across processing subracks 62

5 combined with the redundancy of primary and secondary switching fabrics 42A and 42B (along with redundant links 45A and 45B), yields exceptional fault tolerance and scalability.

Although the controller 46 (or controllers 46) may freely allocate needed processing resources across the various processing subracks 62, a certain degree of

10 resource allocation optimization may offer overall call processing capacity improvements to the BSC 40. For example, in the mixed architecture of Fig. 4, the controller 46 may preferentially allocate needed call processing resources such that, where possible, the call processing flow is constrained to a single processing subrack 62.

By preferentially allocating resources common to a given processing subrack 62,

15 the switching resources local to the processing subrack 62 may be utilized to move call traffic data from specific resources in one resource pool 44 to another. This avoids moving call traffic across the communication links 45 and into another processing subrack 62 via the hub subrack 60. Thus, in the illustrated example, the ATM switching resources on a given processing subrack 62 move call traffic between the different

20 resource types (e.g., SO, SOE, etc.) on that subrack, while the given processing subrack's ATM switching resources combine with the hub subrack's centralized (core) ATM switching resources to move call traffic and control information across processing subracks 62.

It is expected that intra-rack ATM switching rates will exceed 300 Mbps, while

25 inter-rack ATM switching rates will exceed 150 Mbps for each redundant link 45A and

45B. Of course, the ATM switching fabric 42 may be designed to take advantage of higher transport rates as needed.

In the mixed architecture illustrated for the processing subracks 62, each processing subrack 62 includes call processing resources from all the resource pools 44,

5 except for PCN-ET 44E resources—although this latter resource pool 44 may also be distributed across the processing subracks 62. By including the basic complement of call processing resources on each subrack 62, scalability of the BSC 40 is enhanced.

From a system operator's perspective, adding a readily quantifiable amount of call processing capacity to the BSC 40 entails adding one or more processing subracks 62,

10 since each processing subrack 62 substantially provides all of the processing resources 44 needed to complete call routing and processing through the BSC 40.

Fig. 5 is a diagram of an exemplary alternate subrack architecture for the BSC

40. In this implementation, each processing subrack 62 is homogeneous in terms of the processing resources it carries. That is, each processing subrack 62 carries a single

15 type of processing resource. In the illustrated example, the processing subrack 62A carries all of the processing devices comprising the FH-ET resource pool 44A, while processing subrack 62B carries the processing devices providing the SOE resource pool

44B, and so on. Note that one or more of the resource pools 44 may be expanded such that more than one processing subrack 62 is dedicated to that resource pool 44.

20 However, even in this scenario, each processing subrack 62 carries only one type of processing device (e.g., FH-ET, SOE, SE, etc.).

To route a given call through the BSC 40, the controller 46 configures the switching fabric 42 to interconnect the specific processing devices needed for the call from one or more of the resource pools 44 across the various processing subracks 62.

25 The controller preferably configures an ATM connection between specific devices in the various resource pools 44 using the primary switching fabric 42A or, in the presence of

failure, the secondary switching fabric 42B. In some cases, the switching fabric 42 may be configured to operate with mixed portions from both the primary and secondary switching fabrics 42A and 42B, respectively.

Of course, the flexibility afforded by selecting a functional route for calls by simply

5 configuring the switching fabric 42 permits the BSC 40 to be implemented in a variety of other hub and processing subrack configurations. For example, a given processing subrack 62 may be configured to carry only two distinct types of call processing resources (two types of resource pools 44) in a “duet” style architecture. The particular types of processing resources carried by any one processing subrack 62 may be
10 selected to minimize switching traffic through the hub subrack 60, or, for example, to enhance scalability from the perspective of system operators.

Fig. 6 better illustrates the ATM-based switching fabric 42A/B in the rack architecture introduced in Fig. 4. Each processing subrack 62 includes redundant ATM switching interfaces 43A and 43B coupling the processing subrack 62 to the centralized
15 ATM switching interface 47A and 47B carried by the hub subrack 60. Here, the switching resources carrying the “A” designation comprise the primary switching fabric 42A, while those carrying the “B” designation comprise the secondary switching fabric 42B. The combination of the primary and secondary switching resources combine generally to form the redundant, distributed switching fabric 42 discussed herein.

20 Fig. 7 illustrates exemplary flow logic for call processing resource allocation in the context of the mixed rack architecture of Fig. 4. Processing begins with selection of the appropriate BH-ET resources from the BH-ET resource pool 44D (block 200). This amounts to selecting the processing subrack 62 that carries the BH-ET resources supporting the specific RBS 52 involved in the call. In general, each processing subrack
25 62 supports specific ones of the RBSs 52 and selection of the appropriate BH-ET

resources from the BH-ET resource pool 44D for a given call depends on which RBS 52 is supporting the UT 18 involved in the call.

Once the processing subrack 62 with the needed BH-ET resources is selected, the controller 46 assigns remaining call processing resources as needed. In an 5 optimized approach, the controller 46 determines if the initially selected processing subrack 62 has the requisite remaining call resources available for completing the call routing through the BSC 40 (block 202). If so, the controller 46 assigns SE, SOE, and FH-ET resources as needed from the portions of the SE 44C, SOE 44B, and FH-ET 44A resource pools carried on the initially selected processing subrack 62 (block 204). In this 10 manner, the call processing remains local to the processing subrack 62, which helps minimize inter-subrack switch traffic routed through the hub subrack 60. The controller 46 may also assign resources from the PCN-ET resource pool 44E as needed.

A needed resource may not be available on the same processing subrack (block 202), in which case the controller 46 determines if a remaining one of the subracks 62 15 has the needed call processing resources available to complete the call processing flow (block 208). If so, the controller makes the necessary resource assignments to complete the call processing flow on that “other” processing subrack 62 and appropriately configures the switching fabric 42 (block 210). Once the necessary resource assignments are made from the various resource pools (e.g., 44A, 44B, 44C, 44D, and 20 44E) and the switching fabric 42 is appropriately configured, call flow resource assignment operations end for the given call (block 206).

If none of the remaining processing subracks 62 have the needed resources available to complete the call processing flow (block 208), the controller 46 makes resource assignments across the minimum number of processing subracks 62 as 25 needed to complete the call processing flow (block 212). Once the necessary resource assignments are made from the various resource pools (e.g., 44A, 44B, 44C, 44D, and

44E) and the switching fabric 42 is appropriately configured, call flow resource assignment operations end for the given call (block 206).

Note that the ability of the controller 46 to constrain call flow resource assignments so that selected devices from the various resource pools 44 reside on the

5 minimum possible number of processing subracks 62 depends to some extent on whether the MSC 14 allows the BSC 40 to assign FH-ET resources from the FH-ET resource pool 44A as needed. Some types of MSCs 14 make the FH-ET link selection during call setup, rather than allowing the BSC 40 to make the selection. Since each processing subrack 62 carries a specific portion of the FH-ET resource pool 44A, the

10 specific FH-ET link assigned by the MSC 14 determines which processing subrack 62 will be used to interface with the MSC 14 for the given call.

Under these circumstances, a given call has a back-end processing subrack selection constraint imposed by which RBS 52 is involved, and a front-end processing subrack selection constraint imposed by which FT-ET link is assigned by the MSC 14.

15 Even so, the BSC 40 can still optimize call flow routing. For example, the controller 46 can bias call processing resource allocation to either the processing subrack 62 with the MSC-assigned FH-ET resources, or the processing subrack 62 with the specific BH-ET resources interfacing with the required RBS 52.

In general, the BSC 40 preferably works to optimize call flow processing resource allocations made from across the resource pools 44 in support of a given call to minimize the number of processing subracks 62 used to support the call. More generally, the BSC 40 works to minimize inter-subrack switching for any given call flow.

20 In support of this, the BSC 44 may be configured to dynamically re-allocate resource assignments for one or more current calls as resources on specific processing subracks 62 within one or more of the resource pools 44 become available during the course of operation.

The present invention may, of course, be carried out in other specific ways than those herein set forth without departing from the spirit and essential characteristics of the invention. The present embodiments are, therefore, to be considered in all respects as illustrative and not restrictive, and all changes coming within the meaning and

5 equivalency range of the appended claims are intended to be embraced therein.